

Majority IP Split

An AI research engagement is usually written up through its model: the architecture, the ablations, the field metrics. The document that actually governs whether the work is buildable, survivable, and ownable is the one nobody publishes, the signed service agreement. We read that agreement from a three-phase subsurface-AI programme built for a major operator in Oman as a board and its legal counsel would read it, and we argue that its structure is a deliberate risk-transfer machine rather than boilerplate. Five clauses carry the transfer. The obligation is written as an effort, not a warranted result, because on genuinely uncertain research a results warranty is a payout the vendor cannot price. Liability is capped at the value of the phase in which a claim arises and never reaches a completed phase's fees, so exposure is a bounded staircase rather than the whole contract. Payment is due within fourteen days of the university receiving the end-client's funding, a back-to-back gate that keeps the vendor from financing the operator's own funding risk indefinitely while placing the vendor at the tail of the chain. Invoicing runs on a five-month cadence over a two-year term, and either party may exit on sixty days' notice, with the vendor able to leave cleanly after a completed phase. The foreground intellectual property splits fifty-one to the operator and forty-nine to the vendor and its company jointly, the smallest majority that still hands control to the operator while preserving the vendor's right to keep improving its own methods. We show the structure holding under load: when Phase-2 compute ran 2.17x over its budget, the contract's own change mechanism carried it as an additional request of USD 104,300 against a USD 533,000 approved programme rather than an emergency reopening. Three interactive instruments make the argument concrete. The deliverable is a reading of the paperwork that a technical buyer, a research vendor, and a board can each use to decide whether an AI engagement is structured to be signed, survived, and owned, or merely to be demonstrated.

Narendra Patwardhan

Research Collaborator

Tarry Singh

Founder & CEO

AI-RESEARCH-CONTRACT / EFFORT-OBLIGATION / LIABILITY-CAP / FOREGROUND-
IP / IP-SPLIT / FUNDING-GATED-PAYMENT

CONTENTS

01	The five levers that move the risk
02	Why the obligation has to be effort, not a result
03	Liability as a bounded staircase, not a cliff
04	Payment gated to the client's funding
05	Sixty-day exits, and the right to leave cleanly
06	The 51/49 foreground-IP split
07	The structure under load: a 2.17x overrun that did not reopen the contract
08	Reading the paperwork as the three parties
09	What the contract, not the model, decided
10	Limitations
11	References

The write-up of an AI engagement almost always ends where the model does. The detector clears its accuracy bar, the ablations line up, the operator's geoscientists nod at the demonstration, and the account closes on the result. The document that decides whether any of that survives contact with a real organisation is the one that never gets published: the service agreement the two sides actually signed. That agreement is where an operator finds out whether it is buying a capability it will own or renting a dependency it will keep paying for, and where a research vendor finds out whether it has priced a bounded piece of work or written itself an open-ended liability on an uncertain outcome. On the day the model works, those two things look identical.

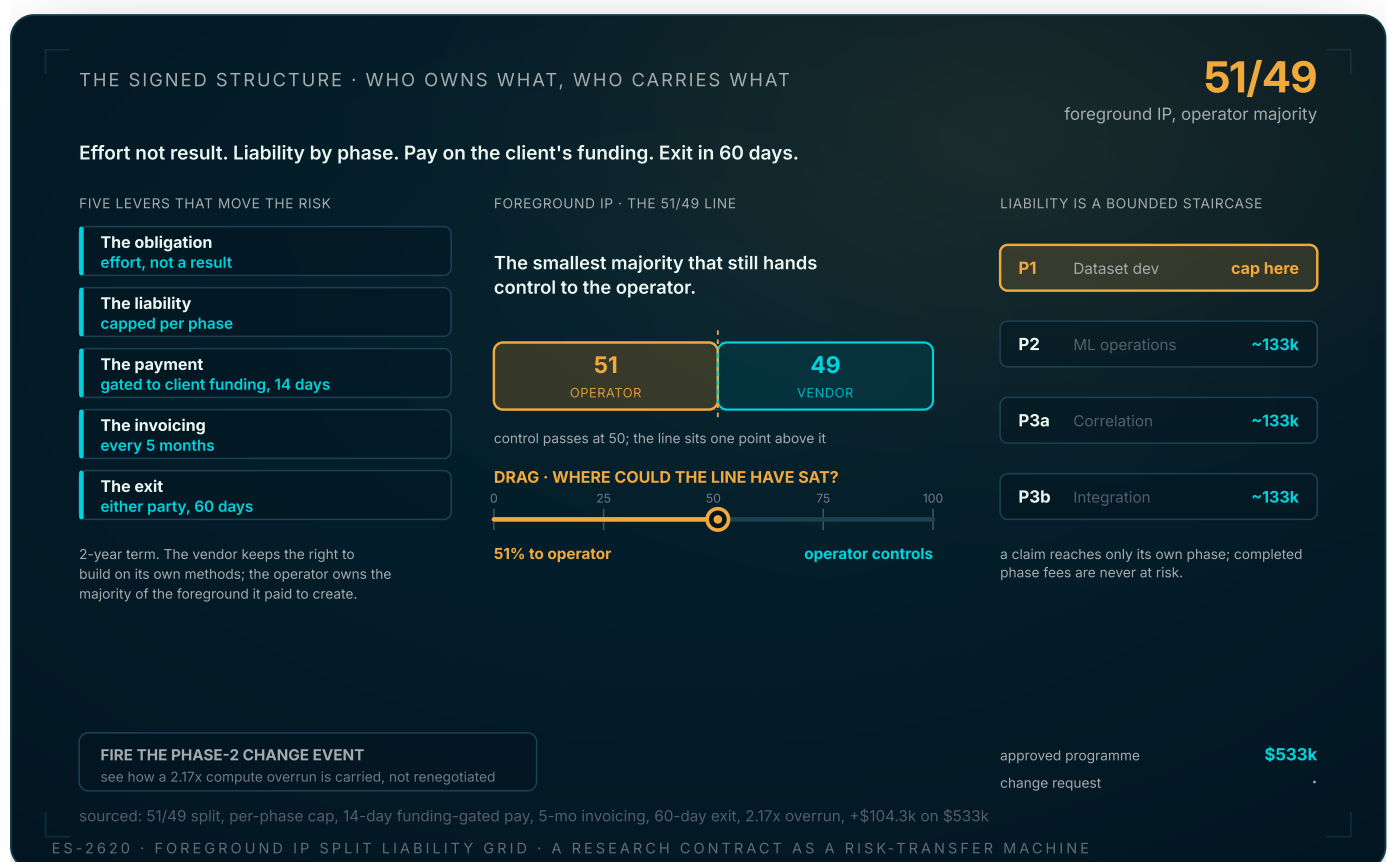
This whitepaper reads that agreement. The programme behind it is a three-phase subsurface-AI build for a major operator in Oman, delivered by an academic partner, Muscat University, and a research vendor working together on borehole-image logs from a fractured carbonate reservoir. The technical story of that programme, the detection-transformer models and the compute economics and the people trained, is told elsewhere and is a pointer here rather than a re-derivation. The subject of this piece is narrower and, we think, more consequential: the shape of the contract, read the way a board and its counsel would read it, as a machine for moving risk between the parties. Five clauses do the work, and they are not incidental legal furniture. They are the reason an uncertain research programme could be signed at all, the reason it survived a serious mid-flight cost overrun without reopening, and the reason it ended with the operator owning the majority of what it paid to create.

The five levers that move the risk

A service agreement for research is, underneath the recitals, a distribution of five risks: the risk that the work does not produce a usable result, the risk that something goes wrong and someone claims damages, the risk that money does not arrive on time, the risk that one party wants out, and the risk over who owns what the work creates. A well-drafted agreement does not pretend those risks away. It assigns each of them, on purpose, to the party best placed to carry it, and it prices the assignment.

The agreement we are reading assigns all five explicitly. The obligation is an effort, not a warranted result. Liability is capped per phase. Payment is gated to the end-client's funding on a fourteen-day window. Invoicing runs on a five-month cadence over a two-year term. Either side can exit on sixty days' notice. And the foreground intellectual property splits

fifty-one to the operator, forty-nine to the vendor and its company jointly. The first instrument lays those five levers side by side and lets you feel where each one sits, because the argument of the whole piece is that they are a system, not a checklist.



A signed research-service contract read as a risk-transfer machine. Left: the five clauses that carry the transfer, an effort obligation rather than a result warranty, liability capped per phase, payment gated to the end-client's funding on a 14-day window, a 5-month invoicing cadence, and a 60-day either-party exit, over a 2-year term. Centre: the foreground-IP division, 51 to the operator and 49 to the vendor and its company jointly; the operator's majority slice is the orange element that argues, and the drag lever sweeps a counterfactual majority so a reader can feel that 51 is the smallest number that still hands control across the 50 threshold, while the vendor keeps the right to build on its own methods. Right: the liability ladder, where a claim reaches only the value of the phase it arises in and never touches a completed phase's fees, so exposure is a bounded staircase rather than the whole contract; the toggle fires the real Phase-2 change event, where compute ran 2.17x over its budget and the contract's own change mechanism carried it as an additional request of USD 104,300 against a USD 533,000 approved programme rather than an emergency reopening. The split, the per-phase cap, the payment and invoicing and notice windows, the term, and the overrun and change-request figures are sourced from the engagement archive; the even per-phase value used to draw the ladder rungs is an illustrative slice of the programme, and none of this is legal advice.


Read left to right, the grid is a single claim: the operator ends up owning the majority of the foreground while carrying almost none of the research risk, and the vendor ends up carrying the research risk in exchange for a bounded downside and the right to keep

improving its own methods. That is a trade both sides can sign. The rest of this piece is an account of why each lever is shaped the way it is, and what breaks if you shape it differently.

Why the obligation has to be effort, not a result

The load-bearing clause is the one that is easiest to miss because it reads like a disclaimer. The vendor is engaged on an effort basis and is, in the contract's own words, not responsible for achieving a specific result, because the engagement is a research project. A buyer new to research procurement often reads that line as the vendor hedging, and pushes to convert it into a performance warranty: pay on the model hitting a stated accuracy, or do not pay. It is an understandable instinct and it is the wrong one, and the reason is not lawyerly caution. It is arithmetic.

A results warranty is, financially, an insurance policy the vendor writes for the buyer. The vendor collects a premium, folded into the price, and pays out when the warranted result fails to materialise. To price that policy honestly, the vendor has to know the probability that the research does not land. On a genuinely uncertain problem, nobody knows that probability, and the vendor's expected payout rises directly with it. The second instrument makes the mechanism visible.



Why an AI research engagement is sold as an effort obligation rather than a results warranty, read as a downside model. The two cards on the left are the same work sold two ways: an effort obligation pays for the work performed with the vendor's downside capped at the phase value and no specific result warranted, while a result warranty pays on a warranted outcome and its downside climbs with the odds the research fails. The plot on the right makes the consequence visible: the teal line is the effort downside, flat at the per-phase cap regardless of how uncertain the outcome is, and the orange curve is the result-warranty downside, rising with research-outcome uncertainty because a warranty pays out exactly when the uncertain research does not land. Drag the uncertainty lever and the marker reads the gap between the two framings, the premium a vendor would have to charge, or refuse, to warrant a result on genuinely uncertain research. The effort framing and the per-phase liability cap are sourced from the engagement archive; the downside-versus-uncertainty curves are an illustrative model rather than measured figures, no money amounts are plotted, and this is not legal or actuarial advice.

Under an effort obligation, the vendor's downside is flat: it is capped at the value of the phase, and it does not move when the research turns out to be harder than hoped, because the vendor was never paid to guarantee the outcome. Under a results warranty, the same downside climbs with the odds the research fails, because that is precisely when the warranty pays out. On a problem where the honest estimate of failure is not near zero, the warranty premium the vendor would have to charge to break even becomes larger than the fee for the work itself, and at that point one of two things happens: the vendor refuses the warranty, or the vendor prices it so high that the buyer refuses the deal. Either way the engagement does not get signed.

The effort framing is what makes it signable. It says, plainly, that the operator is buying competent, well-resourced, honestly-run research effort against a defined scope, and that it is the operator, not the vendor, who carries the irreducible risk that a research question has a disappointing answer. That is the correct assignment. The operator is the party that benefits from the upside if the research lands, the party that chose to pursue an uncertain question, and the party for whom the research is strategic. Asking the vendor to carry the downside of the operator's own research bet, on a fixed fee, is asking the vendor to sell insurance it cannot price, and no competent vendor will. We have written about the sibling of this clause in the year-two services agreement that followed this programme, where the same effort framing does a different job: it makes the vendor safe to pay to work itself out of a job. Here, in year one, it does the more basic thing. It lets the research be bought.

None of this means the effort obligation is a license to underperform. The same agreement carries a scope, a phase structure with defined deliverables, a governance cadence, and the ordinary remedies for a vendor that simply fails to do the work. Effort-not-result draws a line between two very different failures: the failure to do the work competently, which is the vendor's problem, and the failure of a competently-pursued research question to yield the hoped-for answer, which is the operator's. A warranty erases that line and charges the vendor for both. The effort clause preserves it.

Liability as a bounded staircase, not a cliff

The second lever is the liability cap, and its shape matters as much as its existence. Many contracts cap liability at a flat number or at the total contract value. This one caps it at the value of the phase in which a claim arises, and it explicitly protects the fees for already-completed phases from ever being clawed back. The effect is that exposure is a staircase. As the programme advances through its phases, the amount at risk at any moment is the current phase, not the accumulating whole, and the money already earned for finished work is off the table.

That structure is doing something specific for a multi-phase research programme. Research work compounds: Phase 3 depends on the dataset built in Phase 1 and the models trained in Phase 2, and by the time the programme reaches its final phase both parties have sunk substantial effort into the earlier ones. A flat cap at total contract value would mean that a dispute arising in the last phase could, in principle, reach back and put the entire programme's fees at risk, including the fees for phases that were delivered, accepted, and

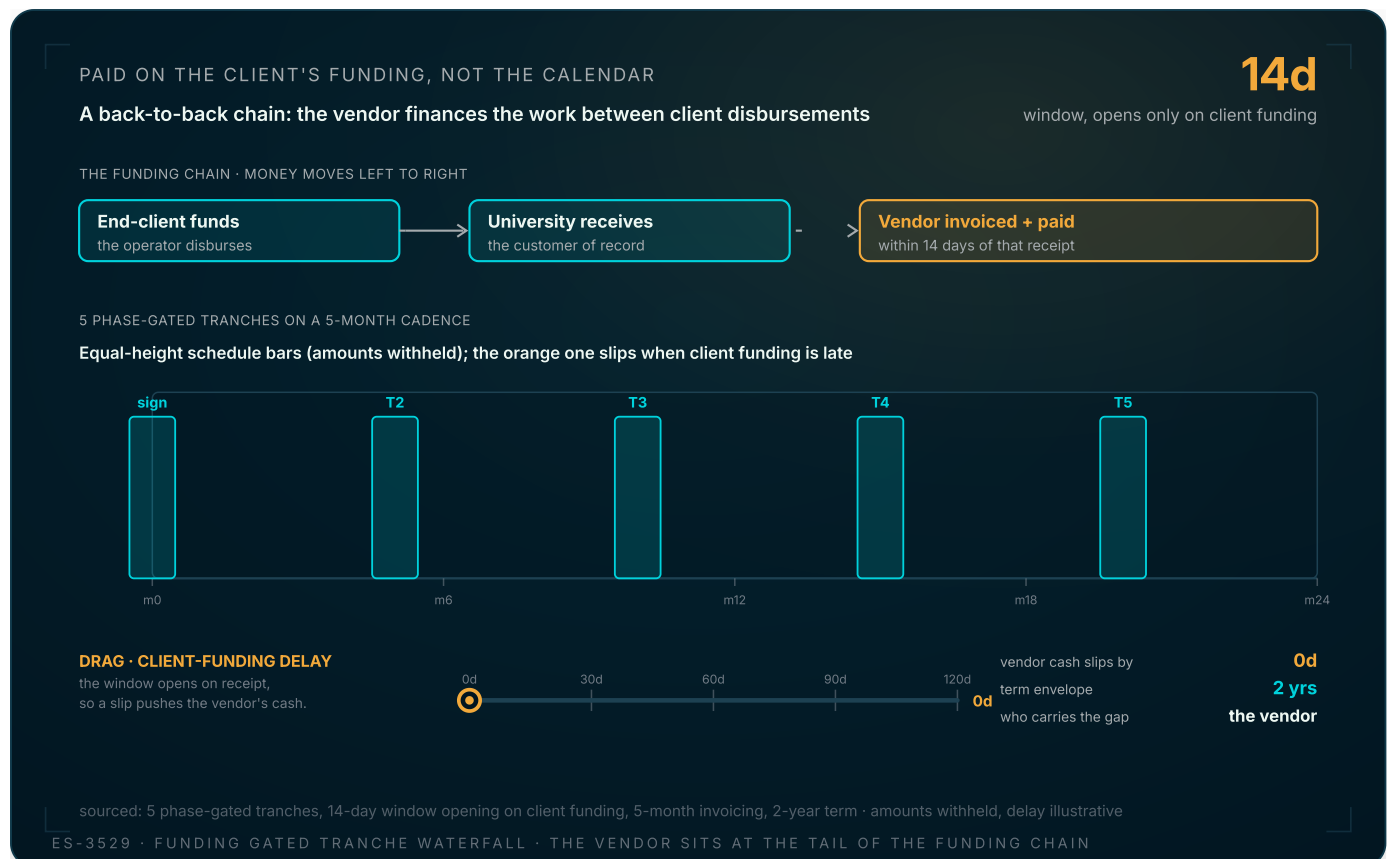
paid months earlier. That is a cliff, and a cliff makes both parties defensive in exactly the wrong way: it discourages the vendor from committing fully to the later, harder phases, because the later the phase the larger the notional exposure. The per-phase staircase removes that distortion. Each phase is a self-contained unit of risk, priced and capped on its own terms, and completing a phase genuinely retires the risk attached to it.

The consequential-damages exclusion sits alongside the cap and completes the picture. The vendor's exposure is bounded not only in amount but in kind: it does not extend to the operator's downstream losses, the drilling decision that went wrong because a dip prediction was off, the production that was deferred. On a programme whose outputs feed structural models that feed field decisions, an uncapped consequential-damages tail would be unquantifiable, and an unquantifiable liability is, again, one the vendor cannot price and therefore will not accept. Excluding it is not the vendor escaping responsibility for its work; it is the two parties agreeing that the operator, who controls the drilling decision and owns the field, is the party that carries the field-scale consequences of that decision.

Payment gated to the client's funding

The third lever is the one that reveals the real shape of the commercial relationship, because it is where the money actually moves. Payment to the vendor is due within fourteen days, but the clock does not start when the vendor invoices. It starts when the university, the customer of record, receives the corresponding funding from the end-client operator. The vendor is paid on the client's funding, not on the calendar.

This is a back-to-back funding gate, and it is a deliberate piece of engineering. The academic partner sits between the operator and the vendor, and without this clause it would be exposed: obliged to pay the vendor on a fixed schedule while itself waiting on disbursements from the operator that might slip. The gate passes that timing risk through, so the university is never financing the operator's funding delays out of its own balance sheet. But the same clause places the vendor at the tail of the chain. When the operator's funding is late, the vendor's cash is late by the same amount, and the vendor is the party carrying the gap. The third instrument traces that chain.



The payment structure read as a financing question. The chain at the top shows the direction of the money: the end-client funds the university, the university receives, and only then does the vendor's 14-day payment window open, so the vendor sits at the tail of a back-to-back funding chain and finances the work between client disbursements. The waterfall plots the five phase-gated tranches, one at signing and four across the two-year term on the 5-month invoicing cadence, drawn at equal height because the absolute amounts are withheld; this is a schedule, not a cash chart. Drag the client-funding delay and the tranche it lands on turns orange, the one element that argues, and slides right by the same amount, leaving a dashed ghost at its scheduled slot so the reader sees the vendor absorbing the gap rather than the client. The tranche count, the 14-day funding-gated window, the 5-month cadence, and the two-year term are sourced from the engagement archive; the tranche amounts are deliberately not plotted, and the client-funding delay swept by the lever is illustrative, not legal advice.

The waterfall shows the five phase-gated tranches across the two-year term, one at signing and four on the five-month invoicing cadence, and it shows what a client-funding delay does to them: the affected tranche slides right by exactly the length of the slip, and the vendor absorbs the wait. We are deliberately not plotting the tranche amounts, which are commercially sensitive; the bars are a schedule, not a cash chart. The point is the structure, not the sums. A vendor signing this clause is agreeing to be the shock absorber for the operator's funding timing, and it should price that agreement, because on a two-year programme a single delayed disbursement can mean months of the vendor financing work already done.

Why would a vendor accept the tail position at all? Because the alternative is worse for everyone. If the university had to guarantee the vendor's payment dates regardless of when the operator funded, it would have to either hold a large working-capital buffer, which raises the price of the whole programme, or refuse the engagement. The back-to-back gate keeps the intermediary solvent and the programme affordable, and it converts the vendor's payment risk into something the vendor can at least see and manage: it is tied to a single, visible event, the operator's disbursement, rather than to the intermediary's general creditworthiness. A vendor that understands where it sits in the chain can manage that. A vendor that does not read the clause discovers its position the first time a tranche is late.

The five-month invoicing cadence is the quiet complement to the funding gate. Invoicing every five months, rather than monthly, matches the billing rhythm to the phase rhythm of the work and keeps the administrative overhead proportionate to a research programme rather than to a staffing contract. It also means the vendor is carrying roughly a phase of work-in-progress between invoices, which is another reason the funding-gate timing matters: the gaps are long enough that a slip is felt.

Sixty-day exits, and the right to leave cleanly

The fourth lever is termination. Either party may end the agreement on sixty days' notice, and the vendor additionally has the right to exit cleanly after a completed phase on short notice. Two years is a long commitment on an uncertain research problem, and the exit rights are what make that length tolerable to both sides.

For the operator, the sixty-day notice is an off-ramp if the research is not going anywhere it values, or if its own priorities shift, without being trapped in the full two-year term. For the vendor, the phase-boundary exit is protection against being dragged into a phase it can see will not work, or into a relationship that has soured, while it is still owed money for completed work, which the liability staircase has already protected. The two exit rights are asymmetric in mechanism but symmetric in intent: each party can leave, neither is a hostage, and the natural exit points align with the phase boundaries where the work has clean seams.

Exit rights on a research contract are not a sign of weak commitment. They are the opposite. A two-year term with no exit forces both parties to negotiate every uncertainty up front, because they cannot revisit it later, which tends to produce a rigid, over-specified

contract that fits the research badly. A term with clean exits lets the parties commit to the current phase wholeheartedly while keeping an honest option to part at the next boundary if the work or the relationship stops making sense. That is the correct posture for research, where the most important information, whether the approach works, arrives during the work rather than before it.

The 51/49 foreground-IP split

The fifth lever is the one that decides who ends up owning the result, and it is the sharpest number in the contract. Foreground intellectual property, meaning the IP created during the engagement rather than the background each party brought to it, splits fifty-one percent to the operator and forty-nine percent to the vendor and its company jointly. Fifty-one is not a round number chosen for symbolism. It is the smallest majority that still hands control to the operator.

Control passes at fifty. A fifty-fifty split is a deadlock: neither party can decide the disposition of the jointly-owned IP without the other, which in practice means it can be neither commercialised nor defended without a negotiation every time. Fifty-one moves the programme one point past deadlock and gives the operator the deciding voice, which is the correct outcome for the party that funded the work and for whom it is strategic. The operator paid for the foreground; the operator should be able to decide what happens to it. The instrument at the top of this piece lets you drag the counterfactual line and watch control flip at fifty, which is the whole reason fifty-one, and not sixty or seventy-five, is the number that got signed. The operator did not need to strip the vendor of a stake to secure control. It needed one point over the line.

The forty-nine that stays with the vendor is not a consolation prize, and it is not really about revenue share. It is about the vendor's right to keep building. A research vendor's entire value is that it accumulates methods across engagements, the architectures, the training recipes, the data-engineering patterns, the hard-won knowledge of what works on this class of problem. If every engagement stripped the vendor of all rights to the foreground, the vendor would be unable to carry forward what it learned, which would make it worse at the next engagement and, paradoxically, worse for the next operator. The forty-nine stake, combined with the ordinary background-IP protections, preserves the vendor's right to keep improving its own methods on the back of what it developed here. The operator gets a

majority stake in this programme's specific foreground and the control that comes with it; the vendor keeps the ability to remain good at its job. Both of those are necessary, and the split is where they are reconciled.

THE SPLIT IN ONE SENTENCE

Fifty-one to the operator is not a revenue formula; it is the minimum majority that gives the funding party control of the foreground while leaving the vendor with a genuine stake and, critically, the freedom to keep building on the methods it brought and improved.

The structure under load: a 2.17x overrun that did not reopen the contract

A contract structure is only worth as much as its behaviour when something goes wrong, and something did. During Phase 2, the compute the work required ran 2.17 times over its budgeted allocation, once the programme split into parallel supervised and unsupervised tracks that roughly doubled the number of concurrent training runs. The mechanics of that overrun, why compute cost is concurrency multiplied by run-length and why the parallel tracks doubled it, are a separate story and we do not re-derive them here. What matters for this piece is what the overrun did to the paperwork.

It did not reopen it. The programme carried the overrun through the contract's own change mechanism, as an additional request of USD 104,300 against a USD 533,000 approved programme, a defined top-up on defined lines rather than an emergency renegotiation of the whole deal. The toggle on the first instrument fires exactly this event, so you can see the change request land against the approved programme without the rest of the structure moving. That is the difference a change mechanism makes. Without one, a serious mid-flight cost surprise forces the parties back to the table to renegotiate the entire agreement, which is slow, adversarial, and often fatal to the relationship, because each side arrives with a year of accumulated grievance. With one, the surprise is handled as a bounded, scoped variation, and the rest of the contract, the effort framing, the liability staircase, the funding gate, the exit rights, the IP split, keeps running untouched.

The change request is also where the effort obligation and the liability cap earn their keep in combination. Because the vendor was engaged on effort and not on a warranted result, a compute overrun is a change in the cost of the effort, not a breach: the vendor did not fail to deliver a guaranteed outcome, the work simply cost more to perform than estimated, and the contract has a clean path for that. Because liability is capped per phase, the overrun conversation is scoped to Phase 2 and does not entangle the completed and paid Phase 1. The structure localises the shock. A different contract, one that warranted a result and capped liability at total value, would have turned the same overrun into a dispute about whether the vendor had breached and about how much of the whole programme's fees were now at risk. This one turned it into a line item.

Reading the paperwork as the three parties

The same structure reads differently depending on which side of the table you sit, and a good agreement is legible to all three.

The technical buyer, the operator's CTO or project director, should read this structure as the thing that lets them procure uncertain research without either overpaying for a warranty they do not need or under-protecting the result they are funding. The buyer's checklist is concrete. Is the obligation effort or result, and if a vendor is offering a results warranty on genuinely uncertain research, what is that warranty actually costing inside the price? Is liability capped in a way that matches the phase structure of the work, or is it a cliff at total value? Does payment timing match the buyer's own funding rhythm, and does the buyer understand that a funding gate places timing risk on the vendor that the vendor is pricing? Are there clean exits at the phase boundaries? And does the IP split give the buyer the control it is paying for, at fifty-one, without stripping the vendor of the ability to stay good at the work?

The research vendor should read it as the structure that lets them take on an uncertain engagement without writing an open-ended liability. The vendor's discipline is to insist on effort-not-result on genuinely uncertain work, to shape the liability cap to the phases so that completing a phase retires its risk, to read the funding gate and price the tail position it puts them in, to keep the phase-boundary exit, and to hold a genuine foreground stake, the forty-nine, for the right to keep building rather than for the revenue. A vendor that gives away all of those to win the deal has won a liability, not a contract.

The board, on either side, should read it as evidence that the engagement is structured to be signed, survived, and owned rather than merely demonstrated. A board does not need to follow the model architecture. It needs to see that the risks have been named and assigned to the parties best placed to carry them, that the downside is bounded and quantifiable, that a mid-flight surprise has a defined path that does not detonate the relationship, and that ownership of the result lands where the strategy requires. The five levers, read together, are that evidence. The accuracy numbers tell the board whether the model works. The contract tells the board whether the programme will hold.

What the contract, not the model, decided

The temptation in writing up an AI programme is to let the model carry the whole story, because the model is where the visible cleverness lives. The model mattered. But the model was buildable because the obligation was effort and not a warranty the vendor could not price; the programme survived a 2.17x compute overrun because the contract had a change mechanism and a per-phase liability staircase that localised the shock; the intermediary stayed solvent because payment was gated to the operator's funding; both parties could commit to a two-year term because either could exit cleanly at a phase boundary; and the operator ended up owning the majority of the foreground it paid for because the split was set at fifty-one, one point past deadlock, without denying the vendor the stake it needed to keep improving its methods.

Every one of those outcomes was a choice made in the paperwork, before a single model was trained. That is the argument. The structure of a research-service agreement is not legal overhead to be minimised on the way to the interesting technical work. It is the mechanism that decides whether the interesting technical work can be bought, can survive its own surprises, and ends up owned by the party that funded it. Read the contract the way you would read the model. It is doing at least as much work.

WHAT THIS WHITEPAPER ARGUES

- 1 A research-service agreement is a risk-transfer machine, not boilerplate: five clauses, effort-not-result, per-phase liability cap, funding-gated payment, 60-day exits, and a 51/49 foreground-IP split, each assign a specific risk to the party best placed to carry it.
- 2 AI R&D must be sold as an effort obligation, not a results warranty, because a warranty on genuinely uncertain research is a payout the vendor cannot price: its expected downside climbs with the odds the research fails, while an effort obligation keeps the downside flat and bounded at the phase value.
- 3 Liability capped at the value of the phase a claim arises in, protecting completed-phase fees, makes exposure a bounded staircase rather than a cliff at total contract value, which is the correct shape for compounding, multi-phase research.
- 4 Payment gated to within 14 days of the end-client funding the intermediary keeps the intermediary solvent and the programme affordable, but places the vendor at the tail of the funding chain, a timing risk the vendor must read and price; invoicing runs on a 5-month cadence over a 2-year term.
- 5 The 51/49 split is the smallest majority that hands the operator control of the foreground while leaving the vendor a genuine stake and the right to keep building on its own methods; the structure held under a real 2.17x Phase-2 compute overrun, carried as a USD 104,300 change request against a USD 533,000 programme rather than an emergency reopening.

Limitations

This is a reading of one signed agreement from one three-phase subsurface-AI programme, and it is a structural reading, not legal advice. The specific numbers, the fifty-one to forty-nine split, the fourteen-day funding-gated window, the five-month invoicing cadence, the sixty-day notice, the two-year term, and the USD 104,300 change request against a USD 533,000 approved programme, are drawn from the engagement's own contract and change-request records; the absolute tranche amounts are commercially sensitive and are deliberately not disclosed. The exposure model in the second instrument is an illustration of how downside scales with outcome uncertainty under the two framings,

not a measured actuarial curve, and no money amounts are plotted on it; the per-phase value used to draw the liability ladder in the first instrument is an even illustrative slice of the programme rather than the true phase-by-phase allocation. Contract structure that is appropriate for an academic-plus-vendor delivery to a national operator under one governing law will not transfer unchanged to a different jurisdiction, a different party structure, or a different risk appetite. Nothing here should be read as a template to be copied; it is an argument about what the levers do, which each reader must apply to their own facts and their own counsel. Finally, the client identity and the specific well and field names are withheld throughout under the engagement's confidentiality terms, and the operator is described only as a major operator in Oman.

References

Muscat University and research vendor, 2021. Service agreement for a machine-learning-for-well-data-analytics research programme (confidential). The signed effort-obligation framing, per-phase liability cap, funding-gated fourteen-day payment window, five-month invoicing cadence, two-year term, sixty-day termination, and fifty-one to forty-nine foreground-IP split read in this piece are drawn from this agreement.

Research vendor and academic partner, 2022. Additional-budget request against the approved programme (confidential). The Phase-2 compute overrun of 2.17x and the additional request of USD 104,300 against the USD 533,000 approved programme are drawn from this record.

Papering an AI Research Engagement: Effort Obligations, Phase-Capped Liability, and a Majority IP Split

Published March 2026. Generated 2026-07-22.

<https://www.earthscan.io/whitepapers/papering-an-ai-research-engagement-effort-obligations-liability-ip-split>

Copyright EarthScan. All rights reserved.